

# **BEACON MONITOR OPERATIONS ON THE DEEP SPACE ONE MISSION**

*E. Jay Wyatt, Henry Hotz, Robert Sherwood,  
John Szijjarto, Miles Sue*

e.j.wyatt@jpl.nasa.gov  
818-354-1414  
818-393-6004 FAX

Jet Propulsion Laboratory  
California Institute of Technology  
Pasadena, California 90019

## **ABSTRACT**

A new approach to mission operations will be flight validated on NASA's New Millennium Program Deep Space One (DS1) mission scheduled to launch in October of 1998. The beacon monitor technology is an operational concept for reducing deep space mission operations cost and decreasing the loading on ground antennas. The technology is required for upcoming NASA missions to Pluto and Europa. With beacon monitoring, the spacecraft will assess its own health and will transmit one of four sub-carrier frequency tones to inform the ground how urgent it is to track the spacecraft for telemetry. If all conditions are nominal, the tone provides periodic assurance to ground personnel that the mission is proceeding as planned without having to receive and analyze downlinked telemetry. If there is a problem, the tone will indicate that tracking is required and the resulting telemetry will contain a concise summary of what has occurred since the last telemetry pass. The primary components of the technology are a tone messaging system, AI-based software for onboard engineering data summarization, a ground visualization system for telemetry summaries, and a ground response system. This paper will describe the operational concept, key hardware and software components, the flight validation approach and the DS1 development status. Applicability to future missions will also be included.

## **1.0 INTRODUCTION**

The operational approach was conceived approximately three years ago in order to lower the cost of the planned NASA mission to Pluto. The New Millennium Program accepted the technology and it was selected for flight validation on the Deep Space One (DS1) mission. Beacon operations on DS1 will be conducted as an experiment and once validated the system will be available for use in DS1 operations. In addition to being baselined for the upcoming missions to Pluto, Europa, and the Sun, several other deep space missions under development are interested in using this approach. The technology is also being funded by the Deep Space Network (DSN) organization as a means for reducing the loading on their large aperture antennas.

## **2.0 DEEP SPACE MISSION OPERATIONS CONCEPT**

The primary objectives of this technology are to lower total mission cost and to decrease the loading on DSN antennas. The fact that NASA full-cost accounting requires that new missions pay for tracking cost is a major motivating factor for finding innovative approaches to operations. The following are major themes in the operational concept:

- Substantially reduce the frequency of telemetry tracking
- Enable the spacecraft to determine the frequency of contact
- Accommodate varying levels of onboard autonomy (beacon monitoring works for missions with high levels of autonomy as well as for traditional mission designs)
- Conduct operations using shared or on-demand operations teams
- Scale frequency of contact (telemetry and beacon tracking) to mission needs/priorities as mission progresses

In traditional mission operations, the spacecraft receives commands from the ground and in turn transmits telemetry in the form of science or engineering data. Perhaps the best way to think of the beacon message is that it is the *spacecraft* sending a command to the ground that instructs the ground personnel how urgent it is to track the spacecraft for telemetry. There are only four such commands. Thinking of beacon operations in this way forces a paradigm shift over the way we traditionally approach operations. Also, it is very important to not think of the tone message as just a little bit of telemetry. If one does this, it is easy to make the argument that a little more telemetry is better. Our approach is one where telemetry is only transmitted when it is necessary for ground personnel to assist the spacecraft or otherwise very infrequently if the spacecraft is fortunate enough to go long periods (a month or so) without requiring ground assistance. When telemetry tracking is necessary the intelligent data summaries that are transmitted contain the most relevant information but also a complete picture of spacecraft activities the last contact. The key challenge here has been developing an architecture that enables the spacecraft to adaptively create summary information to make best use of the available bandwidth as the mission progresses such that all pertinent data is received in one telemetry pass.

### 3.0 DS1 BEACON MESSAGING SYSTEM

The Beacon Messaging System is responsible for the generation, transmission, and detection of the monitoring signals. There are four monitoring signals, each uniquely represents one of the four urgency-based beacon messages. The tone meanings are summarized in Figure 3.1

Tone	Definition
Nominal	Spacecraft is nominal, all functions are performing as expected. No need to downlink engineering telemetry.
Interesting	An interesting and non-urgent event has occurred on the spacecraft. Establish communication with the ground when convenient to obtain data relating to the event. <u>Example</u> : device reset to clear error caused by SEU, other transient events.
Important	The spacecraft needs servicing. Communication with the ground needs to be achieved within a certain time or the spacecraft state could deteriorate and/or critical data could be lost. <u>Examples</u> : solid state memory near full, non-critical hardware failure.
Urgent	Spacecraft emergency. A critical component of the spacecraft has failed. The spacecraft cannot autonomously recover and ground intervention is required immediately. <u>Examples</u> : 1553 bus failure, PDU failure, SRU failure, IPS gimbal stuck.
- No Tone -	Beacon mode is not operating, spacecraft telecom is not Earth-pointed or spacecraft anomaly prohibited tone from being sent.

Figure 3.1 Tone Definitions

The signal structure is shown in Figure 3.2. Each message is represented by a pair of tones centered about the carrier. These tones are generated by phase-modulating the RF carrier by a squarewave subcarrier using 90 degrees modulation angle. The carrier ( $f_c$ ) is completely suppressed. The resulting downlink spectrum consists of tones at odd multiples of the subcarrier frequency above and below the carrier. Four pairs of tones are needed to represent the four possible messages. For the DS1 experiment, the four subcarrier frequencies ( $f_1, f_2, f_3$ , and  $f_4$ ) are 20, 25, 30, and 35 kHz. Different sets of frequency can be used for different missions.




Figure 3.2 Signal Structure

The monitoring system is designed to achieve a low detection threshold. The goal is to reliably detect the monitoring messages with 0 dB-Hz total-received-signal-to-noise-spectral-density ratio ( $P_t/N_o$ ) using 1000 seconds observation time. Future missions are assumed to carry a low-cost auxiliary oscillator as a frequency source, instead of a more expensive, ultra-stable oscillator. The downlink frequency derived from an auxiliary oscillator is not precisely known due to frequency drifts caused by on-board temperature variations, aging, and uncorrected residual Doppler frequency. In addition, the downlink frequency also exhibits short-term drift and phase noise. These factors must be taken into consideration in the design of the monitoring signal detector.

The monitoring signal detector contains four subcarrier detectors (which are also called tone detectors), one for each message or channel as shown in Figure 3.3. Each subcarrier detector is designed to compute the power spectrum of a pair of baseband channels containing the upper and lower first harmonics of that subcarrier. To evaluate the power spectra, the Fast Fourier Transform algorithm (FFT) is employed for computational efficiency. The FFT (coherent) integration time is limited because of oscillator instability. Experimental and theoretical analysis indicates a proper Fourier integration time of approximately 1 second for signals derived from an on-board auxiliary oscillator. Thus, assuming a 1000-second observation interval, 1000 1-second FFTs are performed on successive segments of data, giving 1000 power spectra.

The power spectra obtained from the 1000 FFTs are then summed to improve the signal to noise ratio. Because of the frequency drift, these spectra must be aligned first before summing. This is accomplished by using a simple frequency-drift model (either a linear, piece-wise linear, or quadratic model) with a range of drift rates constrained by a priori knowledge of the maximum possible frequency drift. The aligned and summed spectra provide the necessary statistics upon which detection decision is made.

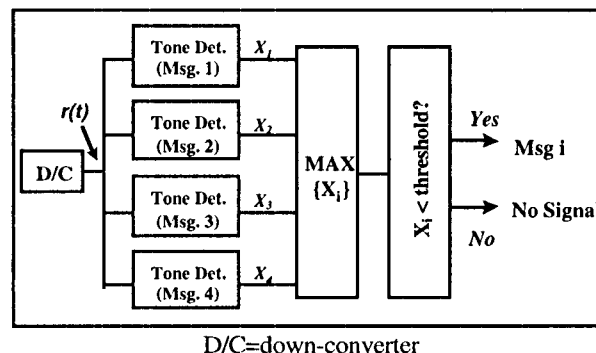


Figure 3.3 Signal Detector and Message Decoder




Figure 3.4a 1-sec Fourier spectra of the four input signals (Provided by G. Lanyi)




Figure 3.4b Fourier spectra of the output after aligning, averaging, and summing 10 FFTs of 1-sec each (Provided by G. Lanyi)

An example of beacon signal detection is shown in Figures 3.4 (a) and (b) using 20 KHz as a signal frequency. Figure 3.4(a) gives the Fourier spectrum of a 1-sec snapshot of the monitoring signal before being processed by the detector, i.e., the spectra of the input signals to the four tone detector. Figure 3.4(b) gives the Fourier spectra of the outputs of the four tone detectors after aligning, summing and averaging over ten FFTs, each of one second duration. The horizontal line is the detection threshold corresponding to a given false alarm probability. As shown in the figure, the aligning and summing process significantly reduces the noise fluctuation and enhances signal detection

The DSN DSS26 (34m) antenna has been assigned as a monitoring station during the mission. The beacon message is first received and decoded by at the Goldstone site and subsequently transmitted to the beacon monitoring team at JPL via a secured link, such as the NASA Science Internet. Next, the beacon message is forwarded to DS1 Mission Operations and other end users, including the Demand Access Scheduler, using email or pagers. It should be noted that the support provided by DSS-26 is experimental in nature and the service is not backed by the guarantees provided to missions using the DSN operational network.

#### **4.0 DS1 ONBOARD ENGINEERING DATA SUMMARIZATION SYSTEM**

Using BMOX tone monitoring instead of traditional routine ground contacts means detailed spacecraft performance history is unavailable to ground personnel. BMOX includes a data summarization component providing this past data in a concise format. Three different types of summarized data are used: overall performance summary, user-defined performance summary, and anomaly summary. Six different telemetry packets have been defined that contain the summary data (see Figure 4.1). The performance summaries are generated at regular intervals and stored in memory until the next telemetry ground contact. They are computed by applying standard functions, such as minimum, maximum, mean, first derivative, and second derivative, to the data. The summarized data are chosen so the spacecraft state can quickly be determined. User-defined summary data are used for obtaining detailed insight into a particular subsystem and are output at the user's discretion. Anomaly summary data (episodes) are created when the raw and summarized data violate pre-defined high and low limits. These limits are determined by the subsystem specialist and stored in a table on-board the spacecraft. The limit tables are based on the current mission activity. The software also has the capability to use AI-based envelope functions instead of traditional alarm limits. Envelope functions are essentially adaptive alarm limits that are learned by training a neural network with nominal engineering data. the neural net can be onboard or on the ground. For DS1, envelope functions are trained on the ground and then uploaded to the spacecraft.

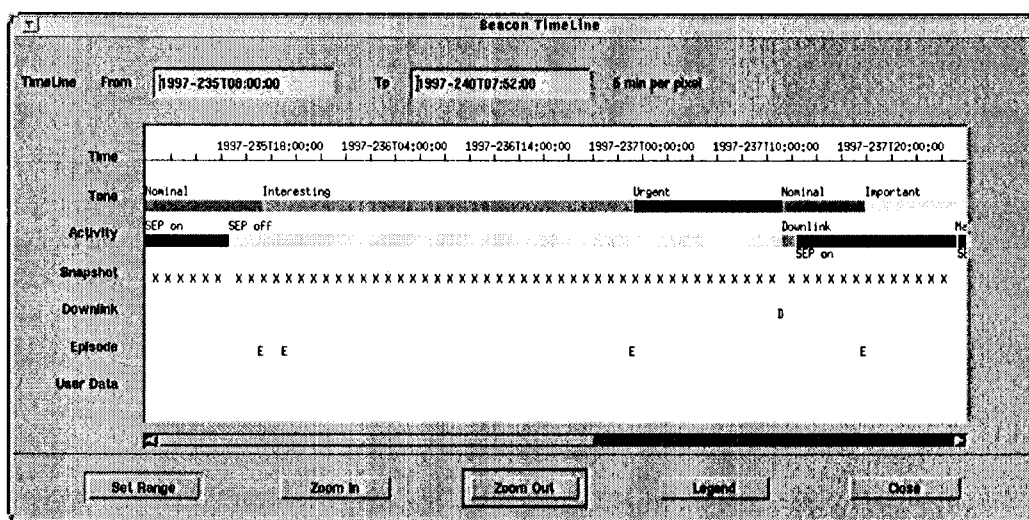
The summarization component combines three functions: data collection and processing, mission activity determination, and episode identification. The data collection subroutine receives engineering data from the engineering telemetry system (EHA) via a function call and applies summary techniques to these data, producing summary measures for downlink to the ground. The mission activity subroutine determines the overall spacecraft mode of operation. This determination is used to choose the appropriate data and limits for a particular episode in the episode subroutine. The mission activity is intended to be exclusive. When a new mission activity starts, the previous mission activity is assumed to have ended. The episode subroutine determines what data is relevant to each particular episode. The episode subroutine combines summary and engineering data received internally from the summarization module with the mission activity received from the activity subroutine and compares the data with mission activity specific alarm limits. For example, Ion Propulsion System (IPS) sensor values may be important while using IPS, but if the spacecraft is in Reaction Control System (RCS) control mode then IPS sensor values could be ignored. In addition, the attitude rate limits might be different during cruise than during a maneuver. As these examples point out, it is necessary to use the mission activities to determine which data to use for episode identification and what are the limits of these data. If the limit is exceeded, the subroutine spawns a new episode and collects past relevant data summarization module. The past data collected will be 1 minute summaries that go back episode length minutes from start of episode. (So a five minute episode would contain summaries starting 5 minutes before the episode to 5 minutes after the episode.) At the end of the episode, the subroutine outputs data to the telemetry subsystem for downlink.

Telemetry Name	Description	Output Frequency
Activity	Documents the changes in mission activity	One packet is produced upon each change
Data Sample	Records a snapshot of every raw and summarized data channel	Regular interval, e.g., 15 minutes
Episode Summary	Records general data about an out-of-limits data condition, known as an "episode."	One per episode
Episode Channel	Records specific data about a single data channel's behavior during an episode.	One or more per episode
Value Summary	Records summary data about a single data channel's behavior since the last downlink.	One for each channel out of limits
User Summary	A user-specified packet containing raw and/or summarized data as specified by the user	Duration user-specified

**Figure 4.1** Summarization Telemetry Packets

The current state-of-the-art in anomaly detection is to use limit-sensing, in which the current sensor value is compared against predetermined high and low "red-lines". Such red-lines are typically constants across many or all mission modes and it is difficult to determine tight limits which will work well throughout the mission. Thus, to avoid frequent false alarms, the red-lines are made imprecise, leading to missed alarms and missed opportunities for early anomaly detection. To compliment the red-line "alarm" type limits we will also employ "warning" type context sensitive envelope functions learned from historical and/or simulated data. These limits are functional values based upon the values of related sensors and other factors, such as the current operational mode of the spacecraft. Although learning precise envelopes can take longer than determining red-lines, initial loose envelopes can be learned quickly. With further training, the bounds can be incrementally tightened, while still retaining a low false alarm rate. Since the learned envelopes are tighter than red-lines, they have a much lower missed alarm rate. Novel training methods are being employed to avoid bounds which cause alarms in nominal training data. Therefore, these envelopes are loose enough to avoid false alarms, provided the training and validation data are representative. In order to learn the envelopes, the ELMER (Envelope Learning and Monitoring via Error Relaxation) algorithm will be used. Pre-flight training will be performed by running ELMER on the ground using historical spacecraft data, examining both anomaly and nominal data sets to determine accurate bounds. For certain phases of the DS-1 beacon monitor experiment, the ground trainer will produce limit functions for uplink.

Tone state and engineering data summaries are displayed on the ground using a special graphical user interface (GUI). The GUI includes a timeline showing all tone changes (detected and



**Figure 4.2** Summary Data Visualization Timeline

telemetry), mission activity changes, snapshot data, downlink summaries, episode data, and user summary data. Figure 3.6 shows the top-level timeline display window. Each of the data types can be displayed in list format or plotted graphically. The environment also includes a strip chart capability and a tool for creating the parameter tables that are uploaded to the spacecraft.

This type of display environment provides a new approach to interacting with telemetry. The basic idea is that the operator should be able to quickly locate important information in the downlink file. If the onboard summarization system is functioning correctly, the most important information will be available at a high enough sample rate to give the operator enough insights into the onboard conditions to perform diagnosis. If the operator searches for data that was never downlinked or is not available at a high enough resolution, then the summarization system has failed. This issue will be addressed in the DS1 flight validation experiments.

## 5.0 TONE RESPONSE SYSTEM

The ground response system processes beacon tone messages, notifying appropriate personnel quickly to facilitate interaction with the spacecraft. The system being developed for demonstration during DS1 is an early prototype that serves the immediate needs on DS1 and also addresses many of the issues associated with developing a system that can serve multiple flight projects. In general when a beacon track occurs the track will be logged and someone will be notified. The form of the notification and its latency depends on the perceived urgency of the event with email assumed for routine events, pager used for significant events requiring prompt attention, and perhaps a synthesized voice call being used for emergencies. Depending on the degree of trust the project has in the notification mechanism it may automatically request antenna time for regular telemetry or emergency tracking. Events are filtered by urgency and type to determine the kind of notification that is required. The notification mechanism should notify specified project positions rather than specific people or addresses. The position should translate to an actual person's contact information based on a duty roster and a personnel database. All notifications must be acknowledged, and the time allowed for the acknowledgment should be configurable on a per-project basis.

The project's interpretation of the signal importance will depend on its operations goals. There are two possible interpretations here. First, the mapping of spacecraft state to urgency of response may evolve as the mission progresses. Early in the prime mission, for example, a device reset may be considered "urgent" because it is wholly unexpected or the consequences are not completely understood. That same event later in the mission, however, may not be considered as urgent and may only trigger the "important" or "interesting" tone. These mappings of spacecraft state to urgency of response can be changed easily by reconfiguring the lookup table. The other interpretation has to do with how each mission defines the latency of response for each tone message. These would vary from mission to mission and may also evolve within a single mission as the operational goals change. Figure 5.1 below illustrates some of the various ways missions may choose to use the beacon signals. The information in this table does not represent actual requirements but rather plausible tone definitions for that mission.

RF state	DS-1	Europa Orbiter (Cruise Ops)	Europa Orbiter Mapping Orbit	Genesis
No Detection	Help	OK	OK	Help
Telemetry	N/A	N/A	N/A	Ignore
Tone 1	OK	OK	OK	Collector Plate 1 OK
Tone 2	Downlink within 2 wks	Dnlink within 1 wk	Dnlink within 3 days (e.g. interesting data)	Collector Plate 2 OK
Tone 3	Dnlink within 2 days	Dnlink within 3 days	Downlink within 1 wk (e.g. buffer filling up)	Collector Plate 3 OK
Tone 4	Urgent, dnlink as soon as possible	Urgent, dnlink as soon as possible	Urgent, dnlink as soon as possible	Downlink within 1 wk

**Figure 5.1** Tone Response - Mission Examples

There are many mission design variables that help determine how a project interprets beacon messages. For example an autonomous spacecraft might be expected to sometimes be out of contact, but not for an extended period of time. Also as long as a finite amount of time is allocated to a Beacon track there will be some uncertainty in the determination of the signal state. Rather than lengthening the required track times it may be desirable to handle the uncertainty directly in the kinds of responses that are taken, in particular a help signal that is not very certain should seek confirmation rather than generating the overhead and expense of a real help signal.

Since the ground response is intended to be automatic whenever an event occurs, it is essential that the system also alert someone when the expected event does not happen. This capability is implemented in a "Backstop Process" which repeatedly checks for various procedural failures. Examples of these failures may be that the automated beacon track did not occur, or the person who was notified did not respond. The actual time allowed for an acknowledgment depends on project requirements. If a notification is not acknowledged, the log of the original notification is so marked and someone else is notified.

## **6.0 FLIGHT VALIDATION OBJECTIVES**

BMOX validates all component technologies and demonstrates an end-to-end operational concept in order to assess operational risks and benefits. There is a set of experiments (or activities) during the mission and also experiment "phases" to describe operational readiness. The experiment phases have been defined as 1) Initial Checkout, 2) Functional Verification, and 3) Operational Effectiveness Assessments. The experiments validate the following components of the technology:

- Engineering summary data generation, visualization, and tone selection
- Tone transmission and detection
- Ground response
- Operational concept assessment and demonstration.

Beacon monitor experiments are suspended during critical spacecraft operation periods including scheduled DSN passes, encounter rehearsals, and actual encounters. The BMOX validation plan during the extended mission includes periods of spacecraft operation solely dependent on beacon monitor tones and beacon-initiated downlinks of summary data.

## **7.0 BENEFIT TO FUTURE MISSIONS**

The primary benefit to future missions is substantial cost savings during mission operations with minimal impact to design-time cost. Although the exact cost of a beacon tone track has not yet been determined, a tone pass is expected to be approximately an order of magnitude cheaper than a telemetry pass. When one considers that routine tracking for current missions often includes daily 8 hour telemetry passes, replacing these with thirty minute tone passes to receive the same assurance of spacecraft health is quite appealing. The tone system has been designed with spacecraft operational constraints in mind. Tones can be transmitted using a low gain antenna during mission periods when telemetry could only be sent with a high gain antenna. Using tones can reduce the amount of consumables required to point the high gain antenna for routine telemetry tracking purposes. This can represent a substantial mass savings during multi-year cruise operations. Missions with solar-electric propulsion or thruster-based attitude control particularly benefit from not having to frequently point the high gain antenna.

## **8.0 REFERENCES**

Wyatt, E. J., et.al (1997) "An Overview of the Beacon Monitor Operations Technology", International Symposium on Artificial Intelligence, Robotics, and Automation in Space, Tokyo, Japan

DeCoste, D. (1997) "Automated Learning and Monitoring of Limit Functions", International Symposium on Artificial Intelligence, Robotics, and Automation in Space, Tokyo, Japan

Sherwood, R., et.al (1997) "Flight Software Implementation of the Beacon Monitor Experiment On the NASA New Millennium Deep Space 1 (DS-1) Mission", Second Intl. Symposium on Reducing the Cost of Spacecraft Ground Systems and Operations, Oxfordshire, UK.

Chien, S., et.al (1997) "Resource Scheduling for a Network of Communications Antennas", Proceedings of the IEEE Aerospace Conference, Aspen, CO.

Staehle, R.L., et.al (1996) "Pluto Express: Advanced Technologies Enable Lower Cost Missions to the Outer Solar System and Beyond" International Low Cost Planetary Missions Conference, Laurel, MD.

Wyatt, E.J. and J.B. Carraway (1995) "Beacon Monitoring Approach to Spacecraft Mission Operations" 1st International Symposium on Reducing the Cost of Spacecraft Ground Systems and Operations, Oxfordshire, UK.